

Hit List

Clear

Generate Collection

Print

Fwd Refs

Bkwd Refs

Generate OACS

Search Results - Record(s) 1 through 4 of 4 returned.

☐ 1. Document ID: US 6330552 B1

L4: Entry 1 of 4

File: USPT

Dec 11, 2001

US-PAT-NO: 6330552

DOCUMENT-IDENTIFIER: US 6330552 B1

TITLE: Database query cost model optimizer

DATE-ISSUED: December 11, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Farrar; Christopher M.	Los Gatos	CA		
Leslie; Harry A.	Los Altos	CA		
Celis; Pedro	Austin	TX		
Shak; Diana L.	San Jose	CA		
Vaishnav; Jay H.	Cupertino	CA		
Skarpelos; Michael J.	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Compaq	Cupertino	CA			02

APPL-NO: 09/ 162638 [PALM]

DATE FILED: September 28, 1998

PARENT-CASE:

CROSS-REFERENCE TO RELATED APPLICATION This application is related to the co-pending U.S. patent application Ser. No. 09/162,245 of Christopher M. Farrar, Harry A. Leslie, Pedro Celis and Diana L. Shak, entitled "Histogram Synthesis Modeler for a Database Query Optimizer," filed Sep. 28, 1998. This patent application is incorporated herein by reference.

INT-CL: [07] G06 F 17/60

US-CL-ISSUED: 705/400; 707/1, 707/2

US-CL-CURRENT: 705/400; 707/1, 707/2

FIELD-OF-SEARCH: 705/400, 707/1, 707/2

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>5301317</u>	April 1994	Lohman et al.	707/2
<u>5367675</u>	November 1994	Cheng et al.	395/600
<u>5546570</u>	August 1996	McPherson, Jr. et al.	395/600
<u>5548755</u>	August 1996	Leung et al.	395/600
<u>5619692</u>	April 1997	Malkemus et al.	395/602
<u>5630120</u>	May 1997	Vachey	395/602
<u>5689696</u>	November 1997	Gibbons et al.	395/601
<u>5689708</u>	November 1997	Regnier et al.	709/229
<u>5761654</u>	June 1998	Tow	707/2
<u>5765146</u>	June 1998	Wolf et al.	707/2
<u>5875445</u>	February 1999	Antonshenkov	707/2
<u>6023695</u>	February 2000	Osborn et al.	707/3
<u>6026391</u>	February 2000	Osbourn et al.	707/2
<u>6212514</u>	April 2001	Eberhard et al.	707/2

FOREIGN PATENT DOCUMENTS

FOREIGN-PAT-NO	PUBN-DATE	COUNTRY	US-CL
11-203288	July 1999	JP	

OTHER PUBLICATIONS

IBM Technical Disclosure Bulletin: "Effective Approach to Query I/O Parallelism Using Sequential Prefetch and Horizontal Data Partitions"; Sep. 1993, vol. 36, No. 9A, pp. 541-544.*

Database Searcher: "Advance cost estimates on DIALOG for search results. (Price Watch)"; Jul.-Aug. 1991, v7, n6, p. 38.*

Spiliopoulou et al: "A cost model for the estimation of query execution time in parallel environment supporting pipelines"; Computers and Artificial Intelligence, 1996, vol. 15, No. 4, pp. 341-368, (Abstract Only).*

Subramanian et al: "Query optimization in multidatabase systems"; Distributed and Parallel Databases, Apr. 1998, vol. 6, No. 2, pp. 183-210, (Abstract Only).*

Article by Clifford A. Lynch entitled "Selectivity Estimation and Query Optimization in Large Databases with Highly Skewed Distributions of Column Values" published by University of California dated 1988 pp. 240-251.

Article by Piatetsky-Shapiro et al. entitled "Accurate Estimation of the Number of Tuples Satisfying a Condition" published by ACM dated 1984 pp. 256-276.

Article by Haas et al. entitled "Sampling-Based Estimation of the Number of Distinct Values of an Attribute" published by Proceedings of the 21st VLDB Conference dated 1995 pp. 311-321.

Article by Yannis E. Ioannidis entitled "University of Serial Histograms" published by Proceedings of the 19th VLDB Conference dated 1993 pp. 256-267.

Article by Mackert et al. entitled "R* Optimizer Validation and Performance Evaluation for Distributed Queries" published by Proceedings of the 12th International Conference on Very Large Databases, dated 1986 pp. 149-159.

Article by Leonard D. Shapiro entitled "Join Processing in Database Systems with Large Main Memories" published by ACM Transactions on Database Systems dated 1986, vol. 11, No. 3, pp. 239-264.

Article by Selinger et al., entitled "Access Path Selection in a Relational Database Management System" published by ACM dated 1979, pp. 23-34.

Article by Sun et al. entitled: "An Instant and Accurate Size Estimation Method for Joins and Selection in a Retrieval-Intensive Environment." SIGMOD, May 1993, pp. 79-88.

Article by Poosala et al. entitled: "Improved Histograms for Selectivity Estimation of Range Predicates" SIGMOD Jun. 1996, pp. 294-305.

ART-UNIT: 211

PRIMARY-EXAMINER: Cosimano; Edward R.

ATTY-AGENT-FIRM: Fenwick & West LLP

ABSTRACT:

A method and apparatus is disclosed for accurately estimating the cost of a database query, including the total computer resources used and estimated elapsed time for the production of a first row and last row of an operator involved in the query and/or the total resources used and elapsed time for returning the overall response to a database query. The method and apparatus accurately accounts for resources used and elapsed time associated with blocking operators, such as sorts and hash joins, which cannot produce a first row until they have completed their operations.

8 Claims, 7 Drawing figures

Full	Title	Citation	Front	Review	Classification	Date	Reference			Claims	KUMC	Draw. Da
------	-------	----------	-------	--------	----------------	------	-----------	--	--	--------	------	----------

☐ 2. Document ID: US 6263345 B1

L4: Entry 2 of 4

File: USPT

Jul 17, 2001

US-PAT-NO: 6263345

DOCUMENT-IDENTIFIER: US 6263345 B1

TITLE: Histogram synthesis modeler for a database query optimizer

DATE-ISSUED: July 17, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Farrar; Christopher M.	Los Gatos	CA		
Leslie; Harry A.	Los Altos	CA		
Celis; Pedro	Austin	TX		
Shak; Diana L.	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Compaq Computers, Corporation	Houston	TX			02

APPL-NO: 09/ 162245 [PALM]

DATE FILED: September 28, 1998

PARENT-CASE:

CROSS-REFERENCE TO RELATED APPLICATION This application is related to the currently co-pending U.S. patent application Ser. No. 09/162,638 of Christopher M. Farrar, Harry A. Leslie, Pedro Celis, Diana L. Shak, Jay H. Vaishnav, and Michael J. Skarpelos, entitled "Database Query Cost Model Optimizer," filed Sep. 28, 1998.

INT-CL: [07] G06 F 17/30

US-CL-ISSUED: 707/104; 707/2, 707/3, 364/468.17, 340/825.44

US-CL-CURRENT: 707/104.1; 700/110, 707/2, 707/3

FIELD-OF-SEARCH: 707/2, 707/3, 707/104, 340/825.44, 364/468.17, 600/407

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>5367675</u>	November 1994	Cheng et al.	395/600
<u>5546570</u>	August 1996	McPherson et al.	395/600
<u>5548755</u>	August 1996	Leung et al.	395/600
<u>5619692</u>	April 1997	Malkemus et al.	395/602
<u>5625815</u>	April 1997	Maier et al.	707/8
<u>5630120</u>	May 1997	Vachey	395/602
<u>5689696</u>	November 1997	Gibbons et al.	395/601
<u>5696686</u>	December 1997	Sanka et al.	364/468.17
<u>5717911</u>	February 1998	Madrid et al.	707/2
<u>5724570</u>	March 1998	Zeller et al.	707/3
<u>5761654</u>	June 1998	Tow	707/2
<u>5803914</u>	September 1998	Ryals et al.	600/407
<u>5819255</u>	October 1998	Celis et al.	707/2
<u>5822747</u>	October 1998	Graefe et al.	707/2
<u>5870752</u>	February 1999	Gibbons et al.	707/102
<u>5942986</u>	August 1999	Shabot et al.	340/825.44
<u>5963957</u>	October 1999	Hoffberg	707/104
<u>6012054</u>	January 2000	Seputis	707/3
<u>6021405</u>	February 2000	Celis et al.	707/2

OTHER PUBLICATIONS

Article by Clifford A. Lynch entitled "Selectivity Estimation and Query Optimization in Large Databases with Highly Skewed Distributions of Column Values" published by University of California, dated 1988 pp. 240-251.

Article by Piatetsky-Shapiro et al. entitled "Accurate Estimation of the Number of Tuples Satisfying a Condition" published by ACM dated, 1984 pp. 256-276.

Article by Haas et al. entitled "Sampling-Based Estimation of the Number of Distinct Values of an Attribute" published by Proceedings of the 21.sup.st VLDB Conference, dated 1995 pp. 311-321.

Article by Yannis E. Ioannidis entitled "Universality of Serial Histograms" published by Proceedings of the 19th VLDB Conference, dated 1993 pp. 256-267.

Article by Mackert et al. Entitled "R* Optimizer Validation and Performance Evaluation for Distributed Queries" published by Proceedings of the 12.sup.th International Conference on Very Large Databases, dated 1986 pp. 149-159.

Article by Leonard D. Shapiro entitled "Join Processing in Database Systems with Large Main Memories" published by ACM Transactions on Database Systems dated 1986, vol. 11, No. 3, pp. 239-264.

Article by Selinger et al., entitled "Access Path Selection in a Relational Database Management System" published by ACM dated 1979, pp. 23-34.

Article by Sun et al. entitled: "An Instant and Accurate Size Estimation Method for Joins and Selection in a Retrieval-Intensive Environment." SIGMOD, May 1993, pp. 79-88.

Article by Poosala et al. entitled: "Improved Histograms for Selectivity Estimation of Range Predicates" SIGMOD Jun. 1996, pp. 294-305.

ART-UNIT: 211

PRIMARY-EXAMINER: Black; Thomas

ASSISTANT-EXAMINER: Rones; Charles L.

ATTY-AGENT-FIRM: Fenwick & West LLP

ABSTRACT:

The invention provides a mechanism for using statistics, in connection with various database query cost modeling techniques, to more accurately estimate the number of rows and UECs that will be produced by relational operators and predicates in database systems. The ability to accurately estimate the number of rows and UECs returned by a relational operator and/or a predicate is fundamental to computing the cost of a query execution plan. This, in turn, drives the optimizer's ability to select the query plan best suited for the desired performance goal. According to the present invention, histogram statistics are synthesized bottom up from the leaf nodes to the root node of a query tree. Given input statistics in the form of histograms for each operand of a relational operator or predicate, the present inventive method and apparatus merge the input statistics in a way that it simulates the effects of the run time operator on the actual data, so as to produce a predicted row count and UEC for each histogram interval representative of the data that actually will be produced by each such operator or predicate in the query tree. A database query optimizer may use these statistics to select and implement an optimal query plan.

14 Claims, 3 Drawing figures

Full	Title	Citation	Front	Review	Classification	Date	Reference			Claims	FIGS	Draw Dg
------	-------	----------	-------	--------	----------------	------	-----------	--	--	--------	------	---------

☐ 3. Document ID: US 6205441 B1

L4: Entry 3 of 4

File: USPT

Mar 20, 2001

US-PAT-NO: 6205441

DOCUMENT-IDENTIFIER: US 6205441 B1

**** See image for Certificate of Correction ****

TITLE: System and method for reducing compile time in a top down rule based system using rule heuristics based upon the predicted resulting data flow

DATE-ISSUED: March 20, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Al-omari; Awny K.	Santa Clara	CA		
Leslie; Harry A.	Los Altos Hills	CA		
Fridrich; Marek J.	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Compaq Computer Corporation	Cupertino	CA			02

APPL-NO: 09/ 283384 [PALM]

DATE FILED: March 31, 1999

PARENT-CASE:

RELATED APPLICATION This application is related to U.S. patent application Ser. No. 09/283,309, filed on Mar. 3, 1999 titled "System and Method for Eliminating Compile Time Explosion in a Top Down Rule Based System Using Selective Sampling," which is incorporated by reference herein in its entirety.

INT-CL: [07] G06 F 17/30

US-CL-ISSUED: 707/2; 707/3, 707/4, 707/5, 707/7

US-CL-CURRENT: 707/2; 707/3, 707/4, 707/5, 707/7

FIELD-OF-SEARCH: 707/2, 707/3, 707/4, 707/5, 707/7

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>4754410</u>	June 1988	Leech et al.	706/45
<u>4964060</u>	October 1990	Hartsog	364/512
<u>5615341</u>	March 1997	Agrawal et al.	705/10
<u>5737728</u>	April 1998	Sisley et al.	705/8
<u>5749069</u>	May 1998	Komori et al.	704/240
<u>5819255</u>	October 1998	Celis et al.	707/2
<u>5822747</u>	October 1998	Graefe et al.	707/2
<u>5855015</u>	December 1998	Shoham	707/5
<u>5914946</u>	June 1999	Avidor et al.	370/336

OTHER PUBLICATIONS

Warshaw et al., "Rule based query optimization, revisited", ACM transaction, pp. 267-275, Nov. 1999.*

Chaudhuri, S. "Optimization of queries with user-defined predicates", ACM transaction on database systems, vol. 24, No. 2, pp. 177-228, Jun. 1999.*

Chakravarthy et al. "Divide and conquer: A basis for augmenting a conventional query optimizer with multiple query-processing capabilities", IEEE, pp. 482-490, Apr. 1991.*

Yoo, Hyuck "Intelligent search method for query optimization by semijoins", IEEE, vol. 1, No. 2, pp. 226-237, Jun. 1989.

ART-UNIT: 271

PRIMARY-EXAMINER: Alam; Hosain T.

ASSISTANT-EXAMINER: Corrielus; Jean M.

ATTY-AGENT-FIRM: Fenwick & West LLP

ABSTRACT:

The present invention applies one or more pruning heuristics to the expression, the binding, and/or the substitute during a database query optimization process. The heuristics identify certain rules that can be eliminated by either not applying the rules and/or not implementing the rules for a given expression and context (if any) based upon one or more flow rates of the expression. The pruning heuristics can eliminate the application of rules based upon the flow rates of the binding or substitute, for example. Examples include (1) not applying (cutting) a MergeJoin rule for a join expression when an inner table is small enough to be stored in a memory space that is allocated for a HashJoin; (2) not applying implementation rules on the substitute of a left-shift rule for an expression if the resulting input data flow rate from the left child of the join is significantly larger in the substitute than in the binding; (3) not applying the join to TSJ (tuple substitute join) rule if the data flow output of the join expression is significantly larger than the data flow input from the inner child of the join expression; or (4) not applying implementation rules on the substitute join expression of a left shift rule if the number of cross products increases and if the data flow rate from the left child is increases.

61 Claims, 35 Drawing figures

Full	Title	Citation	Front	Review	Classification	Date	Reference			Claims	KWAC	Draw De
------	-------	----------	-------	--------	----------------	------	-----------	--	--	--------	------	---------

☐ 4. Document ID: US 5778354 A

L4: Entry 4 of 4

File: USPT

Jul 7, 1998

US-PAT-NO: 5778354

DOCUMENT-IDENTIFIER: US 5778354 A

**** See image for Certificate of Correction ****

TITLE: Database management system with improved indexed accessing

DATE-ISSUED: July 7, 1998

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Leslie; Harry A.	Los Altos Hills	CA		
Birdsall; David W.	Santa Clara	CA		
Jain; Rohit N.	Rochester Hills	MI		
Yaghmai; Hedieh	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Tandem Computers Incorporated	Cupertino	CA			02

APPL-NO: 08/ 481649 [PALM]

DATE FILED: June 7, 1995

INT-CL: [06] G06 F 17/30

US-CL-ISSUED: 707/2; 707/3, 395/182.02, 395/200.03, 395/709, 395/705

US-CL-CURRENT: 707/2; 707/3, 709/201, 714/4, 717/140

FIELD-OF-SEARCH: 395/602, 395/603, 395/182.02, 395/200.03, 395/709, 395/705, 707/2, 707/3

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>4497039</u>	January 1985	Kitakami et al.	364/900
<u>4506326</u>	March 1985	Shaw et al.	364/300
<u>4967341</u>	October 1990	Yamamoto et al.	364/200
<u>5043872</u>	August 1991	Cheng et al.	364/200
<u>5519859</u>	May 1996	Grace	395/600
<u>5551031</u>	August 1996	Cheng et al.	395/600
<u>5557791</u>	September 1996	Cheng et al.	395/600
<u>5559948</u>	September 1996	Bloomfield et al.	395/159
<u>5560006</u>	September 1996	Layden et al.	395/600
<u>5560007</u>	September 1996	Thai	395/600

OTHER PUBLICATIONS

Sartori et al. "Partial Indexing for Nonuniform Data Distributions in Relational DBMS's", IEEE Transaction vol. 6, No. 3, pp. 420-429, Jun. 1994.

IBM Corp., technical Disclosure Bulletin, "Query optimization for relational accessing language", vol. 28, No. 8, pp. 3471-3471, Jan. 1986.

Pang et al., "An efficient semantic query optimization algorithm," Proceeding of the Interl conf on Data Engineering, Kobe, JP, No. 7, IEEE, pp. 326-335, Apr. 1991.

Bayer, R. and Schkolnick, M. "Concurrency of Operations of B-Trees" Acta.Inf. 9 1 (1977).

Bayer, R. and Unterauer, K. "Prefix B-Trees" ACM-TODS, vol. 2, No. 1, Mar. 1977.

Beckmann, N. and Kreigel, H.P. and Schneider, R. and Seeger, "The R*-tree: an Efficient and Robust Access Method for Points and Rectangles". ACM-SIGMOD 1990.

Bentley J. L. "Multidimensional Binary Search Trees Used for Associative Searching" Comm. ACM, vol. 18, No. 9, 1975.

Chang J. M. and Fu K. S. "A Dynamic Clustering Technique for Physical Database Design" ACM-SIGMOD 1980.

Comer, D. "The ubiquitous B-tree" ACM Computing Surveys vol. 11, No. 2, 1979.

Guttman, A. "R-trees: a dynamic index structure for spatial searching" Proceedings ACM-SIGMOD Conf., 1984.

Held, G and Stonebraker, M "B-Trees Re-examined" Electronics Research Laboratory, University of California, 1975.

Knuth, D.E. "The Art of Computer Programming" vol. 3, Sorting and Searching. Addison-Wesley, Reading, Mass., 1973.

Lomet, D.B. and Salzberg, B. "The hB-tree: a Robust Multi-Attribute Indexing

Method" ACM Trans. on Database Systems, vol. 15, No. 4 1989.
Nievergelt, J and Hinterberger, H "The Grid File: An Adaptable, Symmetric Multikey File Structure" Trends in Information Processing Systems, Proc 3rd ECI Conference 1981.
Robinson, J.T. "The K-D-B-Tree: A Search Structure for Large Multi-dimensional Dynamic Indexes" Proc. ACM SIGMOD Conf. 1981.
Rothnie, J. B. and Lozano, T. "Attribute Based File Organization in a Paged Memory Environment" comm ACM, vol. 17, No. 2, Nov. 1974.
Seeger, B. and Kriegel, H. P. "The Buddy-tree: an Efficient and Robust Access Method for Spatial Data Base Systems" Proc. 16th Int. Conf. on Very Large Data Bases, 1990.
Sellis, T. and Rossopoulous, N. and Faloutsos, C. "The R+ Tree: a Dynamic Index for Multidimensional Objects" Proc. 13th Int. Conf. on Very Large Data Bases, 1987.
Snodgrass, R. "Temporal Databases" Computer, pp. 35-42; Computer Sep. 1986.
IBM Corp., IBM Technical Disclosure Bulletin, "Optimization of Like Predicate in an SQL Query", vol. 35, No. 4A, Sep. 1992, pp. 314-317.
IBM Corp., IBM Technical Disclosure Bulletin, "Query Optimization for Relational Accessing Language", vol. 28, No. 8, Jan. 1986, pp. 3471-3472.
Proceedings of the International Conference on Data Engineering, KOBE, JP, Apr. 8-12, 1991, No. 7, IEEE, Hwee Hwa Pang et al., "An Efficient Semantic Query Optimization Algorithm," pp. 326-335.

ART-UNIT: 271

PRIMARY-EXAMINER: Amsbury; Wayne

ASSISTANT-EXAMINER: Corriellus; Jean M.

ATTY-AGENT-FIRM: Graham & James LLP

ABSTRACT:

A database management system (DBMS) provided with a multi-dimensional improved indexed accessing capability using keyed index searching. Individual search keys are constructed from general expression statements created in the DBMS compiler from search queries supplied to the DBMS. Each key column represents another dimension, and both ranges and IN lists can be specified in the search query and used as the predicate values in multiple columns. Missing predicate values in the search query are interpreted as a specification of the minimum and maximum values for the associated search key column. During compile time, the DBMS compiler produces general expressions to be used by the DBMS executor during run time to create the search keys. The DBMS compiler evaluates search queries by associating predicates with clusters and disjunct numbers assigned to each individual disjunct in the search query expression. The DBMS executor uses the general expression from the compiler and eliminates any conflicts among same column predicates, removes redundancies in predicate values and disjuncts and reduces the number of records to be accessed to the minimum required to complete the search query. The individual search keys are generated in the same order as the index to which the search pertains (i.e., increasing or decreasing order).

18 Claims, 11 Drawing figures

Full	Title	Citation	Front	Review	Classification	Date	Reference			Claims	KMC	Draw. De
------	-------	----------	-------	--------	----------------	------	-----------	--	--	--------	-----	----------

Clear

Generate Collection

Print

Fwd Refs

Bkwd Refs

Generate OACS

Term	Documents
"5778354"	1
5778354S	0
((("5778354".PN.) OR 3).USPT.	4
(L3 OR 5778354.PN.).USPT.	4

Display Format:

[Previous Page](#)

[Next Page](#)

[Go to Doc#](#)

Hit List

Clear	Generate Collection	Print	Fwd Refs	Bkwd Refs
Generate OACS				

Search Results - Record(s) 1 through 4 of 4 returned.

☐ 1. Document ID: US 6330552 B1

L4: Entry 1 of 4

File: USPT

Dec 11, 2001

US-PAT-NO: 6330552

DOCUMENT-IDENTIFIER: US 6330552 B1

TITLE: Database query cost model optimizer

DATE-ISSUED: December 11, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Farrar; Christopher M.	Los Gatos	CA		
Leslie; Harry A.	Los Altos	CA		
Celis; Pedro	Austin	TX		
Shak; Diana L.	San Jose	CA		
Vaishnav; Jay H.	Cupertino	CA		
Skarpelos; Michael J.	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Compaq	Cupertino	CA			02

APPL-NO: 09/ 162638 [PALM]

DATE FILED: September 28, 1998

PARENT-CASE:

CROSS-REFERENCE TO RELATED APPLICATION This application is related to the co-pending U.S. patent application Ser. No. 09/162,245 of Christopher M. Farrar, Harry A. Leslie, Pedro Celis and Diana L. Shak, entitled "Histogram Synthesis Modeler for a Database Query Optimizer," filed Sep. 28, 1998. This patent application is incorporated herein by reference.

INT-CL: [07] G06 F 17/60

US-CL-ISSUED: 705/400; 707/1, 707/2

US-CL-CURRENT: 705/400; 707/1, 707/2

FIELD-OF-SEARCH: 705/400, 707/1, 707/2

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>5301317</u>	April 1994	Lohman et al.	707/2
<u>5367675</u>	November 1994	Cheng et al.	395/600
<u>5546570</u>	August 1996	McPherson, Jr. et al.	395/600
<u>5548755</u>	August 1996	Leung et al.	395/600
<u>5619692</u>	April 1997	Malkemus et al.	395/602
<u>5630120</u>	May 1997	Vachey	395/602
<u>5689696</u>	November 1997	Gibbons et al.	395/601
<u>5689708</u>	November 1997	Regnier et al.	709/229
<u>5761654</u>	June 1998	Tow	707/2
<u>5765146</u>	June 1998	Wolf et al.	707/2
<u>5875445</u>	February 1999	Antonshenkov	707/2
<u>6023695</u>	February 2000	Osborn et al.	707/3
<u>6026391</u>	February 2000	Osbourn et al.	707/2
<u>6212514</u>	April 2001	Eberhard et al.	707/2

FOREIGN PATENT DOCUMENTS

FOREIGN-PAT-NO	PUBN-DATE	COUNTRY	US-CL
11-203288	July 1999	JP	

OTHER PUBLICATIONS

IBM Technical Disclosure Bulletin: "Effective Approach to Query I/O Parallelism Using Sequential Prefetch and Horizontal Data Partitions"; Sep. 1993, vol. 36, No. 9A, pp. 541-544.*

Database Searcher: "Advance cost estimates on DIALOG for search results. (Price Watch)"; Jul.-Aug. 1991, v7, n6, p. 38.*

Spiliopoulou et al: "A cost model for the estimation of query execution time in parallel environment supporting pipelines"; Computers and Artificial Intelligence, 1996, vol. 15, No. 4, pp. 341-368, (Abstract Only).*

Subramanian et al: "Query optimization in multidatabase systems"; Distributed and Parallel Databases, Apr. 1998, vol. 6, No. 2, pp. 183-210, (Abstract Only).*

Article by Clifford A. Lynch entitled "Selectivity Estimation and Query Optimization in Large Databases with Highly Skewed Distributions of Column Values" published by University of California dated 1988 pp. 240-251.

Article by Piatetsky-Shapiro et al. entitled "Accurate Estimation of the Number of Tuples Satisfying a Condition" published by ACM dated 1984 pp. 256-276.

Article by Haas et al. entitled "Sampling-Based Estimation of the Number of Distinct Values of an Attribute" published by Proceedings of the 21st VLDB Conference dated 1995 pp. 311-321.

Article by Yannis E. Ioannidis entitled "University of Serial Histograms" published by Proceedings of the 19th VLDB Conference dated 1993 pp. 256-267.

Article by Mackert et al. entitled "R* Optimizer Validation and Performance Evaluation for Distributed Queries" published by Proceedings of the 12th International Conference on Very Large Databases, dated 1986 pp. 149-159.

Article by Leonard D. Shapiro entitled "Join Processing in Database Systems with Large Main Memories" published by ACM Transactions on Database Systems dated 1986, vol. 11, No. 3, pp. 239-264.

Article by Selinger et al., entitled "Access Path Selection in a Relational Database Management System" published by ACM dated 1979, pp. 23-34.

Article by Sun et al. entitled: "An Instant and Accurate Size Estimation Method for Joins and Selection in a Retrieval-Intensive Environment." SIGMOD, May 1993, pp. 79-88.

Article by Poosala et al. entitled: "Improved Histograms for Selectivity Estimation of Range Predicates" SIGMOD Jun. 1996, pp. 294-305.

ART-UNIT: 211

PRIMARY-EXAMINER: Cosimano; Edward R.

ATTY-AGENT-FIRM: Fenwick & West LLP

ABSTRACT:

A method and apparatus is disclosed for accurately estimating the cost of a database query, including the total computer resources used and estimated elapsed time for the production of a first row and last row of an operator involved in the query and/or the total resources used and elapsed time for returning the overall response to a database query. The method and apparatus accurately accounts for resources used and elapsed time associated with blocking operators, such as sorts and hash joins, which cannot produce a first row until they have completed their operations.

8 Claims, 7 Drawing figures

Full	Title	Citation	Front	Review	Classification	Date	Reference			Claims	KWAC	Draw D
------	-------	----------	-------	--------	----------------	------	-----------	--	--	--------	------	--------

☐ 2. Document ID: US 6263345 B1

L4: Entry 2 of 4

File: USPT

Jul 17, 2001

US-PAT-NO: 6263345

DOCUMENT-IDENTIFIER: US 6263345 B1

TITLE: Histogram synthesis modeler for a database query optimizer

DATE-ISSUED: July 17, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Farrar; Christopher M.	Los Gatos	CA		
Leslie; Harry A.	Los Altos	CA		
Celis; Pedro	Austin	TX		
Shak; Diana L.	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Compaq Computers, Corporation	Houston	TX			02

APPL-NO: 09/ 162245 [PALM]

DATE FILED: September 28, 1998

PARENT-CASE:

CROSS-REFERENCE TO RELATED APPLICATION This application is related to the currently co-pending U.S. patent application Ser. No. 09/162,638 of Christopher M. Farrar, Harry A. Leslie, Pedro Celis, Diana L. Shak, Jay H. Vaishnav, and Michael J. Skarpelos, entitled "Database Query Cost Model Optimizer," filed Sep. 28, 1998.

INT-CL: [07] G06 F 17/30

US-CL-ISSUED: 707/104; 707/2, 707/3, 364/468.17, 340/825.44

US-CL-CURRENT: 707/104.1; 700/110, 707/2, 707/3

FIELD-OF-SEARCH: 707/2, 707/3, 707/104, 340/825.44, 364/468.17, 600/407

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>5367675</u>	November 1994	Cheng et al.	395/600
<u>5546570</u>	August 1996	McPherson et al.	395/600
<u>5548755</u>	August 1996	Leung et al.	395/600
<u>5619692</u>	April 1997	Malkemus et al.	395/602
<u>5625815</u>	April 1997	Maier et al.	707/8
<u>5630120</u>	May 1997	Vachey	395/602
<u>5689696</u>	November 1997	Gibbons et al.	395/601
<u>5696686</u>	December 1997	Sanka et al.	364/468.17
<u>5717911</u>	February 1998	Madrid et al.	707/2
<u>5724570</u>	March 1998	Zeller et al.	707/3
<u>5761654</u>	June 1998	Tow	707/2
<u>5803914</u>	September 1998	Ryals et al.	600/407
<u>5819255</u>	October 1998	Celis et al.	707/2
<u>5822747</u>	October 1998	Graefe et al.	707/2
<u>5870752</u>	February 1999	Gibbons et al.	707/102
<u>5942986</u>	August 1999	Shabot et al.	340/825.44
<u>5963957</u>	October 1999	Hoffberg	707/104
<u>6012054</u>	January 2000	Seputis	707/3
<u>6021405</u>	February 2000	Celis et al.	707/2

OTHER PUBLICATIONS

Article by Clifford A. Lynch entitled "Selectivity Estimation and Query Optimization in Large Databases with Highly Skewed Distributions of Column Values" published by University of California, dated 1988 pp. 240-251.

Article by Piatetsky-Shapiro et al. entitled "Accurate Estimation of the Number of Tuples Satisfying a Condition" published by ACM dated, 1984 pp. 256-276.

Article by Haas et al. entitled "Sampling-Based Estimation of the Number of Distinct Values of an Attribute" published by Proceedings of the 21.sup.st VLDB Conference, dated 1995 pp. 311-321.

Article by Yannis E. Ioannidis entitled "Universality of Serial Histograms" published by Proceedings of the 19th VLDB Conference, dated 1993 pp. 256-267.

Article by Mackert et al. Entitled "R* Optimizer Validation and Performance Evaluation for Distributed Queries" published by Proceedings of the 12.sup.th International Conference on Very Large Databases, dated 1986 pp. 149-159.

Article by Leonard D. Shapiro entitled "Join Processing in Database Systems with Large Main Memories" published by ACM Transactions on Database Systems dated 1986, vol. 11, No. 3, pp. 239-264.

Article by Selinger et al., entitled "Access Path Selection in a Relational Database Management System" published by ACM dated 1979, pp. 23-34.

Article by Sun et al. entitled: "An Instant and Accurate Size Estimation Method for Joins and Selection in a Retrieval-Intensive Environment." SIGMOD, May 1993, pp. 79-88.

Article by Poosala et al. entitled: "Improved Histograms for Selectivity Estimation of Range Predicates" SIGMOD Jun. 1996, pp. 294-305.

ART-UNIT: 211

PRIMARY-EXAMINER: Black; Thomas

ASSISTANT-EXAMINER: Rones; Charles L.

ATTY-AGENT-FIRM: Fenwick & West LLP

ABSTRACT:

The invention provides a mechanism for using statistics, in connection with various database query cost modeling techniques, to more accurately estimate the number of rows and UECs that will be produced by relational operators and predicates in database systems. The ability to accurately estimate the number of rows and UECs returned by a relational operator and/or a predicate is fundamental to computing the cost of a query execution plan. This, in turn, drives the optimizer's ability to select the query plan best suited for the desired performance goal. According to the present invention, histogram statistics are synthesized bottom up from the leaf nodes to the root node of a query tree. Given input statistics in the form of histograms for each operand of a relational operator or predicate, the present inventive method and apparatus merge the input statistics in a way that it simulates the effects of the run time operator on the actual data, so as to produce a predicted row count and UEC for each histogram interval representative of the data that actually will be produced by each such operator or predicate in the query tree. A database query optimizer may use these statistics to select and implement an optimal query plan.

14 Claims, 3 Drawing figures

Full	Title	Citation	Front	Review	Classification	Date	Reference			Claims	KWC	Draw D
------	-------	----------	-------	--------	----------------	------	-----------	--	--	--------	-----	--------

☐ 3. Document ID: US 6205441 B1

L4: Entry 3 of 4

File: USPT

Mar 20, 2001

US-PAT-NO: 6205441

DOCUMENT-IDENTIFIER: US 6205441 B1

**** See image for Certificate of Correction ****

TITLE: System and method for reducing compile time in a top down rule based system using rule heuristics based upon the predicted resulting data flow

DATE-ISSUED: March 20, 2001

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Al-omari; Awny K.	Santa Clara	CA		
Leslie; Harry A.	Los Altos Hills	CA		
Fridrich; Marek J.	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Compaq Computer Corporation	Cupertino	CA			02

APPL-NO: 09/ 283384 [PALM]

DATE FILED: March 31, 1999

PARENT-CASE:

RELATED APPLICATION This application is related to U.S. patent application Ser. No. 09/283,309, filed on Mar. 3, 1999 titled "System and Method for Eliminating Compile Time Explosion in a Top Down Rule Based System Using Selective Sampling," which is incorporated by reference herein in its entirety.

INT-CL: [07] G06 F 17/30

US-CL-ISSUED: 707/2; 707/3, 707/4, 707/5, 707/7

US-CL-CURRENT: 707/2; 707/3, 707/4, 707/5, 707/7

FIELD-OF-SEARCH: 707/2, 707/3, 707/4, 707/5, 707/7

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>4754410</u>	June 1988	Leech et al.	706/45
<u>4964060</u>	October 1990	Hartsog	364/512
<u>5615341</u>	March 1997	Agrawal et al.	705/10
<u>5737728</u>	April 1998	Sisley et al.	705/8
<u>5749069</u>	May 1998	Komori et al.	704/240
<u>5819255</u>	October 1998	Celis et al.	707/2
<u>5822747</u>	October 1998	Graefe et al.	707/2
<u>5855015</u>	December 1998	Shoham	707/5
<u>5914946</u>	June 1999	Avidor et al.	370/336

OTHER PUBLICATIONS

Warshaw et al., "Rule based query optimization, revisited", ACM transaction, pp. 267-275, Nov. 1999.*

Chaudhuri, S. "Optimization of queries with user-defined predicates", ACM transaction on database systems, vol. 24, No. 2, pp. 177-228, Jun. 1999.*

Chakravarthy et al. "Divide and conquer: A basis for augmenting a conventional query optimizer with multiple query-processing capabilities", IEEE, pp. 482-490, Apr. 1991.*

Yoo, Hyuck "Intelligent search method for query optimization by semijoins", IEEE, vol. 1, No. 2, pp. 226-237, Jun. 1989.

ART-UNIT: 271

PRIMARY-EXAMINER: Alam; Hosain T.

ASSISTANT-EXAMINER: Corrielus; Jean M.

ATTY-AGENT-FIRM: Fenwick & West LLP

ABSTRACT:

The present invention applies one or more pruning heuristics to the expression, the binding, and/or the substitute during a database query optimization process. The heuristics identify certain rules that can be eliminated by either not applying the rules and/or not implementing the rules for a given expression and context (if any) based upon one or more flow rates of the expression. The pruning heuristics can eliminate the application of rules based upon the flow rates of the binding or substitute, for example. Examples include (1) not applying (cutting) a MergeJoin rule for a join expression when an inner table is small enough to be stored in a memory space that is allocated for a HashJoin; (2) not applying implementation rules on the substitute of a left-shift rule for an expression if the resulting input data flow rate from the left child of the join is significantly larger in the substitute than in the binding; (3) not applying the join to TSJ (tuple substitute join) rule if the data flow output of the join expression is significantly larger than the data flow input from the inner child of the join expression; or (4) not applying implementation rules on the substitute join expression of a left shift rule if the number of cross products increases and if the data flow rate from the left child is increases.

61 Claims, 35 Drawing figures

Full	Title	Citation	Front	Review	Classification	Date	Reference			Claims	KWIC	Draw D
------	-------	----------	-------	--------	----------------	------	-----------	--	--	--------	------	--------

☐ 4. Document ID: US 5778354 A

L4: Entry 4 of 4

File: USPT

Jul 7, 1998

US-PAT-NO: 5778354

DOCUMENT-IDENTIFIER: US 5778354 A

**** See image for Certificate of Correction ****

TITLE: Database management system with improved indexed accessing

DATE-ISSUED: July 7, 1998

INVENTOR-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY
Leslie; Harry A.	Los Altos Hills	CA		
Birdsall; David W.	Santa Clara	CA		
Jain; Rohit N.	Rochester Hills	MI		
Yaghmai; Hedieh	San Jose	CA		

ASSIGNEE-INFORMATION:

NAME	CITY	STATE	ZIP CODE	COUNTRY	TYPE CODE
Tandem Computers Incorporated	Cupertino	CA			02

APPL-NO: 08/ 481649 [PALM]
 DATE FILED: June 7, 1995

INT-CL: [06] G06 F 17/30

US-CL-ISSUED: 707/2; 707/3, 395/182.02, 395/200.03, 395/709, 395/705
 US-CL-CURRENT: 707/2; 707/3, 709/201, 714/4, 717/140

FIELD-OF-SEARCH: 395/602, 395/603, 395/182.02, 395/200.03, 395/709, 395/705, 707/2, 707/3

PRIOR-ART-DISCLOSED:

U.S. PATENT DOCUMENTS

PAT-NO	ISSUE-DATE	PATENTEE-NAME	US-CL
<u>4497039</u>	January 1985	Kitakami et al.	364/900
<u>4506326</u>	March 1985	Shaw et al.	364/300
<u>4967341</u>	October 1990	Yamamoto et al.	364/200
<u>5043872</u>	August 1991	Cheng et al.	364/200
<u>5519859</u>	May 1996	Grace	395/600
<u>5551031</u>	August 1996	Cheng et al.	395/600
<u>5557791</u>	September 1996	Cheng et al.	395/600
<u>5559948</u>	September 1996	Bloomfield et al.	395/159
<u>5560006</u>	September 1996	Layden et al.	395/600
<u>5560007</u>	September 1996	Thai	395/600

OTHER PUBLICATIONS

Sartori et al. "Partial Indexing for Nonuniform Data Distributions in Relational DBMS's", IEEE Transaction vol. 6, No. 3, pp. 420-429, Jun. 1994.

IBM Corp., technical Disclosure Bulletin, "Query optimization for relational accessing language", vol. 28, No. 8, pp. 3471-3471, Jan. 1986.

Pang et al., "An efficient semantic query optimization algorithm," Proceeding of the Interl conf on Data Engineering, Kobe, JP, No. 7, IEEE, pp. 326-335, Apr. 1991.

Bayer, R. and Schkolnick, M. "Concurrency of Operations of B-Trees" Acta.Inf. 9 1 (1977).

Bayer, R. and Unterauer, K. "Prefix B-Trees" ACM-TODS, vol. 2, No. 1, Mar. 1977.

Beckmann, N. and Kreigel, H.P. and Schneider, R. and Seeger, "The R*-tree: an Efficient and Robust Access Method for Points and Rectangles". ACM-SIGMOD 1990.

Bentley J. L. "Multidimensional Binary Search Trees Used for Associative Searching" Comm. ACM, vol. 18, No. 9, 1975.

Chang J. M. and Fu K. S. "A Dynamic Clustering Technique for Physical Database Design" ACM-SIGMOD 1980.

Comer, D. "The ubiquitous B-tree" ACM Computing Surveys vol. 11, No. 2, 1979.

Guttman, A. "R-trees: a dynamic index structure for spatial searching" Proceedings ACM-SIGMOD Conf., 1984.

Held, G and Stonebraker, M "B-Trees Re-examined" Electronics Research Laboratory, University of California, 1975.

Knuth, D.E. "The Art of Computer Programming" vol. 3, Sorting and Searching. Addison-Wesley, Reading, Mass., 1973.

Lomet, D.B. and Salzberg, B. "The hB-tree: a Robust Multi-Attribute Indexing

Method" ACM Trans. on Database Systems, vol. 15, No. 4 1989.

Nievergelt, J and Hinterberger, H "The Grid File: An Adaptable, Symmetric Multikey File Structure" Trends in Information Processing Systems, Proc 3rd ECI Conference 1981.

Robinson, J.T. "The K-D-B-Tree: A Search Structure for Large Multi-dimensional Dynamic Indexes" Proc. ACM SIGMOD Conf. 1981.

Rothnie, J. B. and Lozano, T. "Attribute Based File Organization in a Paged Memory Environment" comm ACM, vol. 17, No. 2, Nov. 1974.

Seeger, B. and Kriegel, H. P. "The Buddy-tree: an Efficient and Robust Access Method for Spatial Data Base Systems" Proc. 16th Int. Conf. on Very Large Data Bases, 1990.

Sellis, T. and Rossopoulous, N. and Faloutsos, C. "The R+ Tree: a Dynamic Index for Multidimensional Objects" Proc. 13th Int. Conf. on Very Large Data Bases, 1987.

Snodgrass, R. "Temporal Databases" Computer, pp. 35-42, Computer Sep. 1986.

IBM Corp., IBM Technical Disclosure Bulletin, "Optimization of Like Predicate in an SQL Query", vol. 35, No. 4A, Sep. 1992, pp. 314-317.

IBM Corp., IBM Technical Disclosure Bulletin, "Query Optimization for Relational Accessing Language", vol. 28, No. 8, Jan. 1986, pp. 3471-3472.

Proceedings of the International Conference on Data Engineering, KOBE, JP, Apr. 8-12, 1991, No. 7, IEEE, Hwee Hwa Pang et al., "An Efficient Semantic Query Optimization Algorithm," pp. 326-335.

ART-UNIT: 271

PRIMARY-EXAMINER: Amsbury; Wayne

ASSISTANT-EXAMINER: Corriellus; Jean M.

ATTY-AGENT-FIRM: Graham & James LLP

ABSTRACT:

A database management system (DBMS) provided with a multi-dimensional improved indexed accessing capability using keyed index searching. Individual search keys are constructed from general expression statements created in the DBMS compiler from search queries supplied to the DBMS. Each key column represents another dimension, and both ranges and IN lists can be specified in the search query and used as the predicate values in multiple columns. Missing predicate values in the search query are interpreted as a specification of the minimum and maximum values for the associated search key column. During compile time, the DBMS compiler produces general expressions to be used by the DBMS executor during run time to create the search keys. The DBMS compiler evaluates search queries by associating predicates with clusters and disjunct numbers assigned to each individual disjunct in the search query expression. The DBMS executor uses the general expression from the compiler and eliminates any conflicts among same column predicates, removes redundancies in predicate values and disjuncts and reduces the number of records to be accessed to the minimum required to complete the search query. The individual search keys are generated in the same order as the index to which the search pertains (i.e., increasing or decreasing order).

18 Claims, 11 Drawing figures

Full	Title	Citation	Front	Review	Classification	Date	Reference			Claims	RMK	Draw D
------	-------	----------	-------	--------	----------------	------	-----------	--	--	--------	-----	--------

Clear	Generate Collection	Print	Fwd Refs	Bkwd Refs	Generate OACS
-------	---------------------	-------	----------	-----------	---------------

Term	Documents
"5778354"	1
5778354S	0
((("5778354".PN.) OR 3).USPT.	4
(L3 OR 5778354.PN.).USPT.	4

Display Format:

[Previous Page](#)

[Next Page](#)

[Go to Doc#](#)